

Application No. 09/784761 (Docket: BAN.0103)
37 CFR 1.111 Amendment dated 07/12/2006
Reply to Office Action of 02/14/2006

RECEIVED
CENTRAL FAX CENTER
JUL 12 2006

AMENDMENTS TO THE SPECIFICATION

Please delete the section entitled "SUMMARY OF THE INVENTION" in its entirety and substitute the following section therefor:

SUMMARY OF THE INVENTION

To address the above-detailed deficiencies, it is an object of the present invention to provide an apparatus that performs the TCP/IP-related processing functions normally attributed to a server.

Accordingly, in the attainment of the aforementioned object, it is a feature of the present invention to provide a TCP-aware target adapter, for accelerating TCP/IP connections between a plurality of clients and a plurality of servers. The plurality of servers are accessed via an Infiniband fabric and the plurality of clients are accessed via a TCP/IP network. The TCP-aware target adapter includes an accelerated connection processor and a target channel adapter. The accelerated connection processor bridges TCP/IP transactions between the plurality of clients and the plurality of servers, where the accelerated connection processor accelerates the TCP/IP connections by ~~intercepting commands provided to the servers' TCP/IP stacks and~~ prescribing Infiniband remote direct memory access operations to retrieve/provide transaction data from/to the plurality of servers. The target channel adapter is coupled to the accelerated connection processor. The target channel adapter supports Infiniband operations with the plurality of servers, and executes the remote direct memory access operations to retrieve/provide the transaction data. The accelerated connection processor has a connection correlator that is configured to associate TCP/IP connection parameters with a target work queue number for each of a plurality of accelerated TCP/IP connections. The TCP/IP connections are accelerated by offloading TCP/IP processing otherwise performed by the plurality of servers to retrieve/provide said transaction data.

An advantage of the present invention is that a server's capacity to perform other processing functions is significantly increased.

Application No. 09/784761 (Docket: BAN.0103)
37 CFR 1.111 Amendment dated 07/12/2006
Reply to Office Action of 02/14/2006

Another object of the present invention is to provide an apparatus in a server that allows TCP/IP transaction data to be transferred to a client machine without requiring that the server perform the processing to decompose the transaction data into packets and to execute TCP/IP transactions to transfer the packets to the client machine.

In another aspect, it is a feature of the present invention to provide an apparatus in a server connected to an Infiniband fabric for implementing accelerated TCP/IP connections between the server and clients. The clients are connected to a TCP/IP network. The apparatus has a connection acceleration driver and a host channel adapter. The connection acceleration driver manages the accelerated TCP/IP connections, ~~where the accelerated TCP/IP connections are established via intercepting commands provided to the server's TCP/IP stack, and~~ where the connection acceleration driver designates memory locations within server memory such that transaction data can be retrieved/provided via Infiniband remote direct memory access operations. The host channel adapter is coupled to the connection acceleration driver. The connection acceleration driver has correlation logic that is configured to associate TCP/IP connection parameters with a host work queue number for each of the accelerated TCP/IP connections. The host channel adapter executes Infiniband operations via the Infiniband fabric, and executes direct memory access functions to retrieve/provide the transaction data responsive to the Infiniband remote direct memory access operations. The accelerated TCP/IP connections offload TCP/IP processing otherwise performed by the server to retrieve/provide said transaction data.

Another advantage of the present invention is that servers no longer need be closely tied to performing protocol-related operations to ensure that data is provided to clients on a network.

Yet another object of the present invention is to provide an apparatus and method for rapidly transferring data from a server to clients connected to a TCP/IP network.

In yet another aspect, it is a feature of the present invention to provide an apparatus within a client-server environment for managing an accelerated TCP/IP connection between a server connected to an Infiniband fabric and a client connected to a TCP/IP

Application No. 09/784761 (Docket: BAN.0103)
 37 CFR 1.111 Amendment dated 07/12/2006
 Reply to Office Action of 02/14/2006

network. The apparatus includes a host driver and a TCP-aware target adapter. The host driver provides a host work ~~queue~~ queue pair through which transaction data corresponding to the accelerated TCP/IP connection is transmitted/received via the Infiniband fabric, ~~where the host driver intercepts a command provided to the server's TCP/IP stack to accelerate the accelerated TCP/IP connection.~~ The TCP-aware target adapter is coupled to the host driver. The TCP-aware target adapter provides a target work ~~queue~~ queue pair corresponding to the host work ~~queue~~ queue pair. The TCP-aware target adapter executes a remote direct memory access operation to receive/transmit the transaction data via the Infiniband fabric. The TCP-aware target adapter includes a connection correlator, for associating TCP/IP connection parameters for the accelerated connection with the a target work queue number corresponding to the target work queue pair. The accelerated TCP/IP connection offloads TCP/IP processing otherwise performed by the server to receive/transmit said transaction data.

In a further aspect, it is a feature of the present invention to provide a method for accelerating TCP/IP connections in a client-server environment having clients that are connected to a TCP/IP network and servers that are connected to an Infiniband fabric. The method includes ~~first intercepting requests to send/receive data provided to the servers' TCP/IP stacks,~~ mapping TCP/IP connection parameters for accelerated connections to corresponding host and target work queue numbers that correspond to host and target work queue pairs, and offloading TCP/IP processing otherwise performed by the servers by executing Infiniband remote direct memory access operations to retrieve/transmit data associated with the accelerated connections from/to memory within the servers.

In yet a further aspect, it is a feature of the present invention to provide a method for offloading server TCP/IP processing in a client-server environment. The method includes bypassing a TCP/IP stack otherwise employed in a server by ~~intercepting send/receive commands provided to the TCP/IP stack and~~ utilizing remote direct memory access operations via an Infiniband fabric to directly access data from/to server memory, where the data is provided to/from a TCP-aware target adapter, the TCP-aware target adapter providing native network ports that connect to clients; and via the TCP-aware

Application No. 09/784761 (Docket: BAN.0103)
37 CFR 1.111 Amendment dated 07/12/2006
Reply to Office Action of 02/14/2006

target adapter, generating native network transactions to transfer the data to/from clients. The utilizing includes associating TCP/IP connection parameters for a particular TCP/IP connection with a work queue number that corresponds to a work queue pair within the TCP-aware target adapter; and issuing remote direct memory access requests to the work queue pair.

In still another aspect, it is a feature of the present invention to provide a TCP-aware target adapter, for accelerating TCP/IP connections between a plurality of clients and a plurality of servers, the plurality of servers being accessed via an Infiniband fabric, the plurality of clients being accessed via a TCP/IP network. The TCP-aware target adapter has an accelerated connection processor and a target channel adapter. The accelerated connection processor bridges TCP/IP transactions between the plurality of clients and the plurality of servers, where the accelerated connection processor accelerates the TCP/IP connections by prescribing remote direct memory access operations to retrieve/provide transaction data from/to the plurality of servers. The accelerated connection processor includes a connection correlator that is configured to associate TCP/IP connection parameters which uniquely identify the TCP/IP connections with corresponding work queue numbers. The target channel adapter is coupled to the accelerated connection processor. The target channel adapter supports Infiniband operations with the plurality of servers, and executes the remote direct memory access operations to retrieve/provide the transaction data, and routes the transaction data to/from the plurality of clients as embedded payloads within Infiniband packets. The TCP/IP connections are accelerated by offloading TCP/IP processing otherwise performed by the plurality of servers to retrieve/provide said transaction data.

In yet another aspect, it is a feature of the present invention to provide an Infiniband-to-native protocol translation apparatus, for routing TCP/IP transactions between a plurality of clients and a plurality of Infiniband devices. The plurality of Infiniband devices are accessed via an Infiniband fabric and the plurality of clients are accessed via a TCP/IP network. The Infiniband-to-native protocol translation apparatus has an unaccelerated connection processor and a target channel adapter. The unaccelerated connection processor bridges the TCP/IP transactions between the plurality of clients and the

Application No. 09/784761 (Docket: BAN.0103)
37 CFR 1.111 Amendment dated 07/12/2006
Reply to Office Action of 02/14/2006

plurality of Infiniband devices by ~~intercepting send/receive commands provided to corresponding TCP/IP stacks within the plurality of servers and encapsulating/stripping~~ the TCP transactions within/from Infiniband raw packets. The unaccelerated connection processor includes an unaccelerated connection correlator. The unaccelerated connection correlator maps native addresses to/from Infiniband local identifiers and work queue numbers. The target channel adapter is coupled to the unaccelerated connection processor. The target channel adapter receives/transmits said Infiniband raw packets from/to the plurality of Infiniband devices.

Another advantage of the present invention is that the number of servers within a data center can be increased over the Infiniband fabric without impacting the techniques employed to interface the servers to a client TCP/IP network.